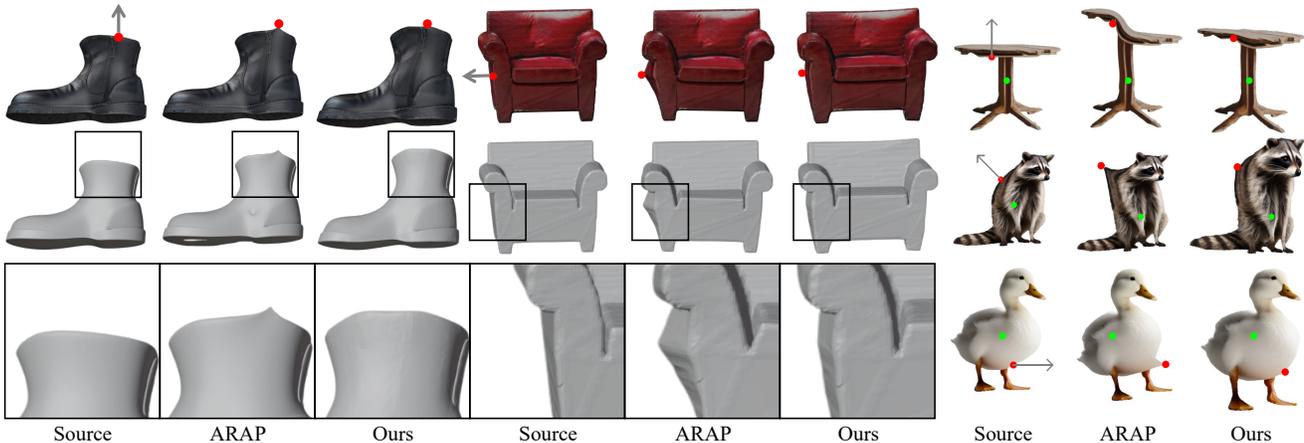


# As-Plausible-As-Possible: Plausibility-Aware Mesh Deformation Using 2D Diffusion Priors

Seungwoo Yoo<sup>\*1</sup> Kunho Kim<sup>\*1</sup> Vladimir G. Kim<sup>2</sup> Minhyuk Sung<sup>1</sup>  
<sup>1</sup>KAIST <sup>2</sup>Adobe Research



**Figure 1.** APAP, our novel shape deformation method, enables plausibility-aware mesh deformation and preservation of fine details of the original mesh offering an interface that alters geometry by directly displacing a handle (*red*) along a direction (*gray*), fixing an anchor vertex (*green*). Using a diffusion prior results in smoother geometry around the armchair handle, as seen in the example (middle column).

## Abstract

We present *As-Plausible-as-Possible (APAP) mesh deformation technique that leverages 2D diffusion priors to preserve the plausibility of a mesh under user-controlled deformation. Our framework uses per-face Jacobians to represent mesh deformations, where mesh vertex coordinates are computed via a differentiable Poisson Solve. The deformed mesh is rendered, and the resulting 2D image is used in the Score Distillation Sampling (SDS) process, which enables extracting meaningful plausibility priors from a pretrained 2D diffusion model. To better preserve the identity of the edited mesh, we fine-tune our 2D diffusion model with LoRA. Gradients extracted by SDS and a user-prescribed handle displacement are then back-propagated to the per-face Jacobians, and we use iterative gradient descent to compute the final deformation that balances between the user edit and the output plausibility. We evaluate our method with 2D and 3D meshes and demonstrate qualitative and quantitative improvements when using plausibility priors over geometry-preservation or distortion-minimization priors used by previous techniques. Our project page is at: <https://as-plausible-as-possible.github.io/>*

<sup>\*</sup>Equal contribution.

## 1. Introduction

For 2D and 3D content, mesh is the most prevalent representation, thanks to its efficiency in storage, simplicity in rendering and also compatibility in common graphics pipelines, versatility in diverse applications such as design, physical simulation, and 3D printing, and flexibility in terms of decomposing geometry and appearance information, with widespread adoption in the industry.

For the creation of 2D and 3D meshes, recent breakthroughs in generative models [18, 23, 24, 29, 30, 32, 36, 38] have demonstrated significant advances. These breakthroughs enable users to easily generate content from a text prompt [23, 24, 30, 36, 38], or from photos [25, 30]. However, visual content creation typically involves numerous editing processes, deforming the content to satisfy users’ desires through interactions such as mouse clicks and drags. Facilitating such interactive editing has remained relatively underexplored in the context of recent generative techniques.

Mesh deformation is a subject that has been researched for decades in computer graphics. Over time, researchers have established well-defined methodologies, characterizing mesh deformation as an optimization problem that aims to preserve specific geometric properties, such as the Mesh Laplacian [20, 21, 34], local rigidity [8, 33], and

mesh surface Jacobians [1, 6], while satisfying given constraints. To facilitate user interaction, these methodologies have been extended to introduce specific user-interactive deformation handles, such as keypoints [10, 16, 37], cage mesh [12, 14, 15, 19, 39, 43], and skeleton [2, 41, 42], with the blending functions defined based on the preservation of geometric properties.

Despite the widespread use of classical mesh deformation methods, they often fail to meet users’ needs because they do not incorporate the perceptual plausibility of the outputs. For example, as illustrated in Fig. 1, when a user intends to drag a point on the top of a table image, the classical deformation technique may introduce unnatural bending instead of lifting the tabletop. This limitation arises because deformation techniques solely based on geometric properties do not incorporate such semantic and perceptual priors, resulting in the mesh editing process becoming more tedious and time-consuming.

Recent learning-based mesh deformation techniques [1, 12, 16, 22, 35, 41, 43] have attempted to address this problem in a data-driven way. However, they are also limited by relying on the existence of certain variations in the training data. Even recent large-scale 3D datasets [3–5, 40] have not reached the scale that covers all possible visual content users might intend to create.

To this end, we introduce our novel mesh deformation framework, dubbed APAP (As-Plausible-As-Possible), which exploits 2D image priors from a diffusion model pretrained on an Internet-scale image dataset to enhance the plausibility of deformed 2D and 3D meshes while preserving the geometric priors of the given shape. Recently, score distillation sampling (SDS) [24] has demonstrated great success in generating plausible 2D and 3D content, such as NeRF [13, 17, 44] and vector images [9, 11], using the distilled 2D image priors from a diffusion model. We incorporate these diffusion-model-based 2D priors into the optimization-based deformation framework, achieving the best synergy between geometry-based optimization and distilled-prior-based optimization.

## 2. Method

We present **APAP**, a novel handle-based mesh deformation framework capable of producing visually plausible deformations of either 2D or 3D triangular meshes. To achieve this goal, we integrate powerful 2D diffusion priors into a learnable Jacobian field representation of shapes.

We emphasize that leveraging 2D priors, such as latent diffusion models (LDMs) [26] trained on large-scale datasets [28], for shape deformation poses challenges that require meticulous design choices. The following sections will delve into the details of shape representation (Sec. 2.1) and diffusion prior (Sec. 2.2), offering a rationale for the design decisions underpinning our framework (Sec. 2.3).

### 2.1. Representing Shapes as Jacobian Fields

Let  $\mathcal{M}_0 = (\mathbf{V}_0, \mathbf{F}_0)$  denote a source mesh to be deformed, represented by vertices  $\mathbf{V}_0 \in \mathbb{R}^{V \times 3}$  and faces  $\mathbf{F}_0 \in \mathbb{R}^{F \times 3}$ . Users are allowed to select a set of vertices used as movable handles designated by an indicator matrix  $\mathbf{K}_h \in \{0, 1\}^{V_h \times V}$ . We also require users to select a set of anchors, represented as another indicator matrix  $\mathbf{K}_a \in \{0, 1\}^{V_a \times V}$ , to avoid trivial solutions (i.e., global translations). Then, the handle and anchor vertices become  $\mathbf{V}_h = \mathbf{K}_h \mathbf{V}_0$  and  $\mathbf{V}_a = \mathbf{K}_a \mathbf{V}_0$ .

Our framework also expects a set of vectors  $\mathbf{D}_h \in \mathbb{R}^{V_h \times 3}$  that indicate the directions along which the handles will be displaced. Furthermore, we let  $\mathbf{T}_h = \mathbf{V}_h + \mathbf{D}_h$  and  $\mathbf{T}_a = \mathbf{V}_a$  denote the target positions of the user-specified handles and anchors, respectively.

In this work, we employ a Jacobian field  $\mathbf{J}_0 = \{\mathbf{J}_{0,f} | f \in \mathbf{F}_0\}$ , a dual representation of  $\mathcal{M}_0$ , defined as a set of per-face Jacobians  $\mathbf{J}_{0,f} \in \mathbb{R}^{3 \times 3}$  where

$$\mathbf{J}_{0,f} = \nabla_f \mathbf{V}_0, \quad (1)$$

and  $\nabla_f$  is the gradient operator of triangle  $f$ .

Conversely, we compute a set of *deformed* vertices  $\mathbf{V}^*$  from a given Jacobian field  $\mathbf{J}$  by solving a Poisson’s equation

$$\mathbf{V}^* = \arg \min_{\mathbf{V}} \|\mathbf{L}\mathbf{V} - \nabla^T \mathcal{A}\mathbf{J}\|^2, \quad (2)$$

where  $\nabla$  is a stack of per-face gradient operators,  $\mathcal{A} \in \mathbb{R}^{3F \times 3F}$  is the mass matrix and  $\mathbf{L} \in \mathbb{R}^{V \times V}$  is the cotangent Laplacian of  $\mathcal{M}_0$ , respectively. Since  $\mathbf{L}$  is rank-deficient, the solution of Eqn. 2 cannot be uniquely determined unless we impose constraints. We thus consider a constrained optimization problem

$$\mathbf{V}^* = \arg \min_{\mathbf{V}} \|\mathbf{L}\mathbf{V} - \nabla^T \mathcal{A}\mathbf{J}\|^2 + \lambda \|\mathbf{K}_a \mathbf{V} - \mathbf{T}_a\|^2, \quad (3)$$

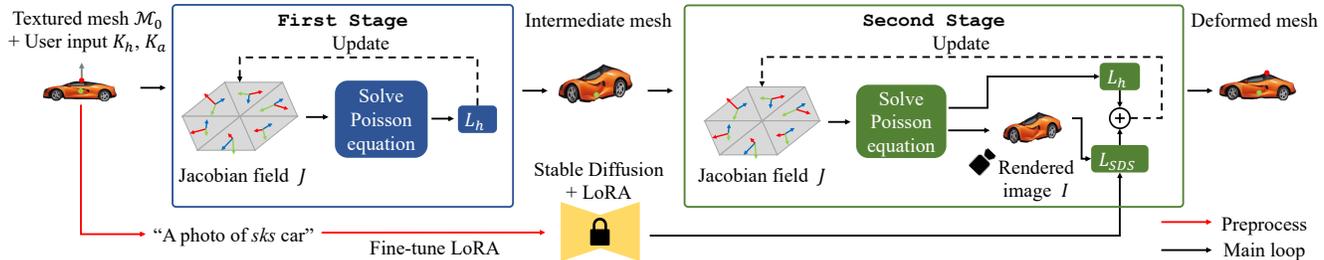
where  $\lambda \in \mathbb{R}^+$  is a weight for the constraint term. Note that we solve Eqn. 3 with the user-specified anchors as constraints to determine  $\mathbf{V}^*$ .

Taking the derivative with respect to  $\mathbf{V}$ , the problem in Eqn. 3 turns into a system of equations

$$(\mathbf{L}^T \mathbf{L} + \lambda \mathbf{K}_a^T \mathbf{K}_a) \mathbf{V} = \mathbf{L}^T \nabla^T \mathcal{A}\mathbf{J} + \lambda \mathbf{K}_a^T \mathbf{T}_a, \quad (4)$$

which can be efficiently solved using a differentiable solver [1] implementing Cholesky decomposition.

We let  $g$  denote a functional representing the aforementioned differentiable solver for notational convenience,  $\mathbf{V}^* = g(\mathbf{J}, \mathbf{K}_a, \mathbf{T}_a)$ . Since  $g$  is differentiable, we can deform  $\mathcal{M}_0$  by propagating upstream gradients from various loss functions to the underlying parameterization  $\mathbf{J}$ . For instance, one may impose a *soft* constraint on the locations of



**Figure 2. The overview of APAP.** APAP parameterizes a triangular mesh as a per-face Jacobian field that can be updated via gradient-descent. Given a textured mesh and user inputs specifying the handle(s) and anchor(s), our framework initializes a Jacobian field as a trainable parameter. During the first stage, the Jacobian field is updated via iterative optimization of  $\mathcal{L}_h$ , a soft constraint that initially deforms the shape according to the user’s instruction. In the following stage, the mesh is rendered using a differentiable renderer  $\mathcal{R}$  and the rendered image is provided as an input to a diffusion prior finetuned with LoRA [7] that computes the SDS loss  $\mathcal{L}_{\text{SDS}}$ . The joint optimization of  $\mathcal{L}_h$  and  $\mathcal{L}_{\text{SDS}}$  improves the visual plausibility of the mesh while conforming to the given edit instruction.

selected handles during optimization with the objective of the form:

$$\mathcal{L}_h = \|\mathbf{K}_h \mathbf{V}^* - \mathbf{T}_h\|^2. \quad (5)$$

We will discuss how such a soft constraint can be blended into our framework in Sec. 2.3. Next, we describe how to incorporate a pretrained diffusion model as a prior for visual plausibility.

## 2.2. Score Distillation for Shape Deformation

While traditional mesh deformation techniques make variations that match the given *geometric* constraints, their lack of consideration on *visual plausibility* results in unrealistic shapes. Motivated by recent success in text-to-3D literature, we harness a powerful 2D diffusion prior [26] in our framework as a critic that directs deformation by scoring the realism of the current shape.

Specifically, we distill its prior knowledge via Score Distillation Sampling (SDS) [24]. Let  $\mathbf{J}$  denote the current Jacobian field and  $\mathbf{V}^*$  be the set of vertices computed from  $\mathbf{J}$  following the procedure described in Sec. 2.1.

We render  $\mathcal{M}^* = (\mathbf{V}^*, \mathbf{F})$  from a viewpoint defined by camera extrinsic parameters  $\mathbf{C}$  using a differentiable renderer  $\mathcal{R}$ , producing an image  $\mathcal{I} = \mathcal{R}(\mathcal{M}^*, \mathbf{C})$ . The diffusion prior  $\hat{\epsilon}_\phi$  then rates the realism of  $\mathcal{I}$ , producing a gradient

$$\nabla_{\mathbf{J}} \mathcal{L}_{\text{SDS}}(\phi, \mathcal{I}) = \mathbb{E}_{t, \epsilon} \left[ w(t) (\hat{\epsilon}_\phi(\mathbf{z}_t; y, t) - \epsilon) \frac{\partial \mathcal{I}}{\partial \mathbf{J}} \right], \quad (6)$$

where  $t \sim \mathcal{U}(0, 1)$ ,  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ , and  $\mathbf{z}_t$  is a noisy latent embedding of  $\mathcal{I}$ . The propagated gradient alters the geometry of  $\mathcal{M}$  by modifying  $\mathbf{J}$ .

To increase the instance-awareness of the diffusion model, we follow recent work [27, 31] on personalized image editing and finetune the model using LoRA [7]. In particular, we first render  $\mathcal{M}$  from  $n$  different viewpoints to

obtain a set  $\mathcal{I} = \{\mathcal{I}_1, \dots, \mathcal{I}_n\}$  of training images and inject additional parameters to the model, resulting in an expanded set of network parameters  $\phi'$ . The parameters are then optimized with a denoising loss [26]

$$\mathcal{L} = \mathbb{E}_{t, \epsilon, \mathbf{z}} [\|\hat{\epsilon}_{\phi'}(\mathbf{z}_t; y, t) - \epsilon\|^2], \quad (7)$$

where  $\mathbf{z}_t$  denotes a latent of a training image perturbed with noise at timestep  $t$ .

The finetuned diffusion prior, together with a learnable Jacobian field representation of the source mesh  $\mathcal{M}_0$ , comprises the proposed framework described in the following section.

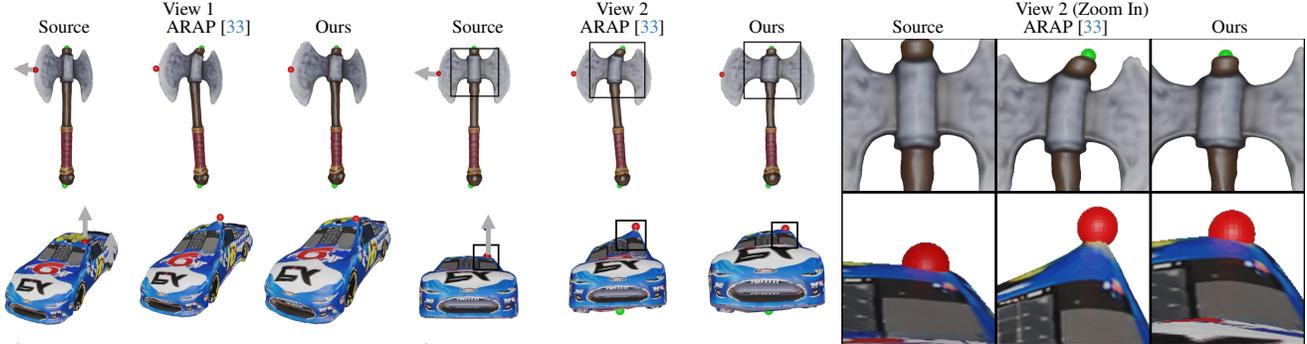
## 2.3. As-Plausible-As-Possible (APAP)

**APAP** tackles the problem of plausibility-aware shape deformation by harmonizing the best of both worlds: a learnable shape representation founded on classical geometry processing, robust to noisy gradients, and a powerful 2D diffusion prior finetuned with the image(s) of the source mesh for better instance-awareness.

We provide an overview of the proposed pipeline in Fig. 2. The detailed algorithm is presented in the **supplementary material**. We delve into details in the following.

Provided with a textured mesh  $\mathcal{M}_0$ , handles  $\mathbf{K}_h$ , anchors  $\mathbf{K}_a$ , as well as their target positions  $\mathbf{T}_h$  and  $\mathbf{T}_a$  as inputs, **APAP** yields a plausible deformation  $\mathcal{M}$  of  $\mathcal{M}_0$  that conforms to the given handle-target constraints. Before deforming  $\mathcal{M}_0$ , we render  $\mathcal{M}_0$  from a single view in the case of 2D meshes and four canonical views (i.e., front, back, left, and right) for 3D meshes and use the images to finetune Stable Diffusion [26] by optimizing LoRA [7] parameters injected to the model (the red line in Fig. 2). Simultaneously, **APAP** computes the Jacobian field  $\mathbf{J}_0$  of the input mesh  $\mathcal{M}_0$  and initializes it as a trainable parameter  $\mathbf{J}$ .

**APAP** deforms the input mesh through two stages. In the *FirstStage*, it first deforms the input mesh according to instructions from users without taking visual plausibility



**Figure 3. Qualitative results from 3D shape deformation.** We visualize the source shapes and their deformations made using ARAP [33] and ours by following the instructions each of which specifies a handle (*red*), an edit direction denoted with an arrow (*gray*), and an anchor (*green*). We showcase the rendered images captured from two different viewpoints, as well as one zoom-in view highlighting local details.

into account. The subsequent *SecondStage* integrates a 2D diffusion prior into the optimization loop, simultaneously enforcing user constraints and visual plausibility.

At every iteration of the *FirstStage* illustrated as the *blue* box in Fig. 2, we compute the vertex positions  $\mathbf{V}^*$  corresponding to the current Jacobian field  $\mathbf{J}$  by solving Eqn. 3 using the anchors specified by  $\mathbf{K}_a$  as hard constraints. Then, we compute the soft constraint  $\mathcal{L}_h$  defined as Eqn. 5 that drags a set of handle vertices  $\mathbf{K}_h \mathbf{V}^*$  toward the corresponding targets  $\mathbf{T}_h$ . The interleaving of differentiable Poisson solve and optimization of  $\mathcal{L}_h$  via gradient-descent is repeated for  $M$  iterations. This progressively updates  $\mathbf{J}$ , treated as a learnable black box in our framework, deforming  $\mathcal{M}_0$ . Consequently, the edited mesh  $\mathcal{M}^* = (\mathbf{J}, \mathbf{F}_0)$  follows user constraints at the cost of the degraded plausibility, mitigated in the following stage through the incorporation of a diffusion prior.

The result of *FirstStage* then serves as an initialization for the *SecondStage*, illustrated as the *green* box in Fig. 2 guided by plausibility constraint  $\mathcal{L}_{\text{SDS}}$ . Unlike the *FirstStage* where the update of  $\mathbf{J}$  was purely driven by the geometric constraint  $\mathcal{L}_h$ , we aim to steer the optimization based on the visual plausibility of the current mesh  $\mathcal{M}^*$ . To achieve this, we render  $\mathcal{M}^*$  using a differentiable renderer  $\mathcal{R}$  using the same viewpoint(s) from which the training image(s) for finetuning was rendered. When deforming 3D meshes, we randomly sample one viewpoint at each iteration. The rendered image  $\mathcal{I}$  is used to evaluate  $\mathcal{L}_{\text{SDS}}$  which is optimized jointly with  $\mathcal{L}_h$  for  $N$  iterations. The combination of geometric and plausibility constraints improves the visual plausibility of the output while encouraging it to conform to the given constraints.

We note that the iterative approach in the *FirstStage* leads to better results than alternative update strategies such as deforming the source mesh  $\mathcal{M}_0$  by minimizing ARAP energy [33] or, solving Eqn. 3 using both  $\mathbf{K}_h$  and  $\mathbf{K}_a$  as hard constraints. In the **supplementary material**, we show that both methods produce distortions that cannot be cor-

rected by the diffusion prior in the subsequent stage. Specifically, directly solving Eqn. 3 using all available constraints only yields the least squares solution  $\mathbf{V}^*$  without updating the underlying Jacobians  $\mathbf{J}$ , resulting in the aforementioned distortions.

### 3. Experiments

We evaluate **APAP** in downstream applications involving manipulation of 3D and 2D meshes. While we focus on 3D mesh manipulation, we summarize experimental details, more 3D deformation examples, and additional results using 2D meshes in the **supplementary material**.

#### 3.1. 3D Shape Deformation

**Qualitative Results.** We showcase examples of 3D shape deformation where each deformation is specified by a handle (*red*), an edit direction (*gray*), and an anchor (*green*). As shown in Fig. 3, **APAP** is capable of manipulating 3D shapes to improve visual plausibility which is not achievable by solely relying on geometric prior such as ARAP [33]. For instance, given a user input that drags a handle on one blade of an axe (the first row) along an arrow, **APAP** simultaneously expands both blades of the axe whereas ARAP [33] produces distortions near the head. Similar examples that demonstrate symmetry-awareness of **APAP** can be found in other cases such as a car (the second row), and an owl (the sixth row) where a user lifts only one side of the shape upward and the symmetry is recovered by **APAP** which cannot be achieved by ARAP [33]. Also, note that **APAP** is capable of making a smooth articulation at the leg of the wolf (the fourth row) by adjusting the overall posture in comparison to ARAP which creates an excess bending.

### 4. Conclusion

We presented **APAP**, a novel deformation framework that tackles the problem of plausibility-aware shape deformation while offering intuitive controls over a wide range of shapes represented as triangular meshes.

## References

- [1] Noam Aigerman, Kunal Gupta, Vladimir G Kim, Siddhartha Chaudhuri, Jun Saito, and Thibault Groueix. Neural Jacobian Fields: Learning Intrinsic Mappings of Arbitrary Meshes. *ACM TOG*, 2022. 2
- [2] Ilya Baran and Jovan Popović. Automatic Rigging and Animation of 3D Characters. *ACM TOG*, 2007. 2
- [3] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. ShapeNet: An Information-Rich 3D Model Repository. *arXiv preprint arXiv:1512.03012*, 2015. 2
- [4] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-XL: A Universe of 10M+ 3D Objects. In *NeurIPS*, 2023.
- [5] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A Universe of Annotated 3D Objects. In *CVPR*, 2023. 2
- [6] William Gao, Noam Aigerman, Groueix Thibault, Vladimir Kim, and Rana Hanocka. TextDeformer: Geometry Manipulation using Text Guidance. *ACM TOG*, 2023. 2
- [7] Yelong Hu, Edward J. and Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*, 2022. 3
- [8] Takeo Igarashi, Tomer Moscovich, and John F. Hughes. As-Rigid-as-Possible Shape Manipulation. *ACM TOG*, 2005. 1
- [9] Shir Iluz, Yael Vinker, Amir Hertz, Daniel Berio, Daniel Cohen-Or, and Ariel Shamir. Word-As-Image for Semantic Typography. *ACM TOG*, 2023. 2
- [10] Alec Jacobson, Ilya Baran, Jovan Popović, and Olga Sorkine. Bounded Biharmonic Weights for Real-Time Deformation. *ACM TOG*, 2011. 2
- [11] Ajay Jain, Amber Xie, and Pieter Abbeel. VectorFusion: Text-to-SVG by Abstracting Pixel-Based Diffusion Models. In *CVPR*, 2023. 2
- [12] Tomas Jakab, Richard Tucker, Ameesh Makadia, Jiajun Wu, Noah Snavely, and Angjoo Kanazawa. KeypointDeformer: Unsupervised 3D Keypoint Discovery for Shape Control. In *CVPR*, 2020. 2
- [13] Jong Chul Ye Jangho Park, Gihyun Kwon. ED-NeRF: Efficient Text-Guided Editing of 3D Scene using Latent Space NeRF. *arXiv*, 2023. 2
- [14] Pushkar Joshi, Mark Meyer, Tony DeRose, Brian Green, and Tom Sanocki. Harmonic Coordinates for Character Articulation. *ACM TOG*, 2007. 2
- [15] Tao Ju, Scott Schaefer, and Joe Warren. Mean Value Coordinates for Closed Triangular Meshes. *ACM TOG*, 2005. 2
- [16] Kunho Kim, Mikaela Angelina Uy, Despoina Paschalidou, Alec Jacobson, Leonidas J. Guibas, and Minhyuk Sung. OptCtrlPoints: Finding the Optimal Control Points for Biharmonic 3D Shape Deformation. *Computer Graphics Forum*, 2023. 2
- [17] Subin Kim, Kyungmin Lee, June Suk Choi, Jongheon Jeong, Kihyuk Sohn, and Jinwoo Shin. Collaborative Score Distillation for Consistent Visual Synthesis. In *NeurIPS*, 2023. 2
- [18] Juil Koo, Seungwoo Yoo, Minh Hieu Nguyen, and Minhyuk Sung. SALAD: Part-Level Latent Diffusion for 3D Shape Generation and Manipulation. In *ICCV*, 2023. 1
- [19] Yaron Lipman, David Levin, and Daniel Cohen-Or. Green Coordinates. *ACM TOG*, 2008. 2
- [20] Yaron Lipman, Olga Sorkine, Daniel Cohen-Or, David Levin, Christian Rössl, and Hans-Peter Seidel. Differential Coordinates for Interactive Mesh Editing. In *Proceedings of Shape Modeling International*, pages 181–190, 2004. 1
- [21] Yaron Lipman, Olga Sorkine, David Levin, and Daniel Cohen-Or. Linear Rotation-Invariant Coordinates for Meshes. *ACM TOG*, 2005. 1
- [22] Minghua Liu, Minhyuk Sung, Radomir Mech, and Hao Su. DeepMetaHandles: Learning Deformation Meta-Handles of 3D Meshes with Biharmonic Coordinates. In *CVPR*, 2021. 2
- [23] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. SyncDreamer: Learning to Generate Multiview-consistent Images from a Single-view Image. *arXiv*, 2023. 1
- [24] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D Diffusion. In *ICLR*, 2023. 1, 2, 3
- [25] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Ben Mildenhall, Nataniel Ruiz, Shiran Zada, Kfir Aberman, Michael Rubenstein, Jonathan Barron, Yuanzhen Li, and Varun Jampani. DreamBooth3D: Subject-Driven Text-to-3D Generation. In *ICCV*, 2023. 1
- [26] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*, 2022. 2, 3
- [27] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. DreamBooth: Fine Tuning Text-to-image Diffusion Models for Subject-Driven Generation. In *CVPR*, 2023. 3
- [28] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5B: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022. 2
- [29] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep Marching Tetrahedra: a Hybrid Representation for High-Resolution 3D Shape Synthesis. In *NeurIPS*, 2021. 1
- [30] Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. MVDream: Multi-view Diffusion for 3D Generation. *arXiv*, 2023. 1
- [31] Yujun Shi, Chuhui Xue, Jiachun Pan, Wenqing Zhang, Vincent YF Tan, and Song Bai. DragDiffusion: Harnessing Diffusion Models for Interactive Point-based Image Editing. *arXiv*, 2023. 3
- [32] J. Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Ankner, Jiajun Wu, and Gordon Wetzstein. 3D Neural Field Generation using Triplane Diffusion. In *CVPR*, 2023. 1

- [33] Olga Sorkine and Marc Alexa. As-Rigid-As-Possible Surface Modeling. *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, pages 109–116, 2007. [1](#), [4](#)
- [34] Olga Sorkine, Daniel Cohen-Or, Yaron Lipman, Marc Alexa, Christian Rössl, and Hans-Peter Seidel. Laplacian Surface Editing. *Proceedings of the EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, pages 179–188, 2004. [1](#)
- [35] Jiapeng Tang, Markhasin Lev, Wang Bi, Thies Justus, and Matthias Nießner. Neural Shape Deformation Priors. In *NeurIPS*, 2022. [2](#)
- [36] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. DreamGaussian: Generative Gaussian Splatting for Efficient 3D Content Creation. *arXiv*, 2023. [1](#)
- [37] Yu Wang, Alec Jacobson, Jernej Barbič, and Ladislav Kavan. Linear Subspace Design for Real-Time Shape Deformation. *ACM TOG*, 2015. [2](#)
- [38] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In *NeurIPS*, 2023. [1](#)
- [39] Ofir Weber, Mirela Ben-Chen, and Craig Gotsman. Complex barycentric coordinates with applications to planar shape deformation. *Computer Graphics Forum*, 2009. [2](#)
- [40] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, Dahua Lin, and Ziwei Liu. OmniObject3D: Large-Vocabulary 3D Object Dataset for Realistic Perception, Reconstruction and Generation. In *CVPR*, 2023. [2](#)
- [41] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. RigNet: Neural Rigging for Articulated Characters. *ACM TOG*, 2020. [2](#)
- [42] Zhan Xu, Yang Zhou, Li Yi, and Evangelos Kalogerakis. Morig: Motion-Aware Rigging of Character Meshes from Point Clouds. In *SIGGRAPH ASIA*, 2022. [2](#)
- [43] Wang Yifan, Noam Aigerman, Vladimir G. Kim, Chaudhuri Siddhartha, and Olga Sorkine. Neural Cages for Detail-Preserving 3D Deformations. In *CVPR*, 2020. [2](#)
- [44] Jingyu Zhuang, Chen Wang, Lingjie Liu, Liang Lin, and Guanbin Li. DreamEditor: Text-Driven 3D Scene Editing with Neural Fields. *ACM TOG*, 2023. [2](#)